

Automatic Tuning of Local Memory Use on GPGPUs

Tianyi David Han Tarek S. Abdelrahman

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering
University of Toronto
Toronto, Ontario M5S 3G4, Canada
{han,tsa}@eecg.toronto.edu

Abstract

The use of local memory is important to improve the performance of OpenCL programs. However, its use may not always benefit performance, depending on various application characteristics, and there is no simple heuristic for deciding when to use it. We develop a machine learning model to decide if the optimization is beneficial or not. We train the model with millions of synthetic benchmarks and show that it can predict if the optimization should be applied for a single array, in both synthetic and real benchmarks, with high accuracy.

1. Introduction

The last few years have seen considerable growth in the use of Graphics Processing Units (GPUs) to accelerate general purpose applications. Today, the OpenCL standard [14] is used to express applications in which computationally intensive segments, or *kernels* are launched for execution on the GPU. However, to realize the performance benefits of GPUs, programmers must *optimize* their kernels to better exploit the underlying GPU architecture [11].

One important optimization is the use of local memory – a small user-managed on-chip storage, which can improve performance by an order of magnitude. However, the optimization is not always beneficial, depending on 1) the amount of data reuse and the degree of memory non-coalescing in the kernel, and 2) the instruction overhead and the amount of drop in parallelism the optimization introduces. The extent to which these factors influence the optimization’s benefit is often not clear. There is no simple heuristic for deciding whether or not the optimization should be applied.

We explore the use of machine learning to auto-tune the local memory optimization. We build a model that predicts the benefit of caching a region of an array in local memory, based on the performance of a set of training kernels with and without the optimization. We then apply this model to a new kernel to decide if the optimization should be applied.

A unique aspect of our work is the use of many synthetically generated kernels for model training. We believe that machine learning, particularly on a high-dimensional fea-

ture space, demands a large training set which is difficult to assemble from real-world benchmarks. The use of synthetic benchmarks allows us to build a robust model fully exploiting the power of machine learning. These synthetic kernels capture common data access patterns in the domains of dense linear algebra and structured grids.

Performance data of a large number of synthetic kernel instances (with and without using local memory) on an NVIDIA GPU shows that the optimization brings a wide range of kernel speedup (from $0.03\times$ to $49.6\times$). A *Random Forest* [2] model trained on a random 10% of the data can predict if the optimization is beneficial, on the remaining data, with nearly 95% accuracy. This model also achieves an average of nearly 95% accuracy on eight real-world kernels.

To our best knowledge, we believe that this is the first work that: 1) builds an accurate model using machine learning to auto-tune the local memory optimization, and 2) uses a large number of synthetic benchmarks for model training.

2. Background

Local memory is a software-managed cache shared by all workitems within a workgroup. Despite being small (in kilobytes), its use can significantly improve kernel performance, because it has close-to-register access latency [11].

Local memory improves performance for two reasons. First, caching data in local memory exploits data locality in the kernel and can reduce the number of transactions reaching the GPU DRAM. Data locality in a kernel (executed by many threads) can be classified as *temporal* or *spatial*, and *intra-thread* or *inter-thread*. Local memory is commonly used to exploit all four categories of data locality except intra-thread temporal locality, where the GPU compilers are able to put the data in thread-private registers. The benefit of using local memory increases with the amount of data reuse.

Even in the absence of data reuse, local memory can improve kernel performance by transforming non-coalesced memory accesses into coalesced ones [15]. A common scenario is when each workitem performs some sort of row-wise reduction, forcing workitems to access a *column* of a 2D array at the same time. This results in totally non-coalesced accesses that come with a high performance

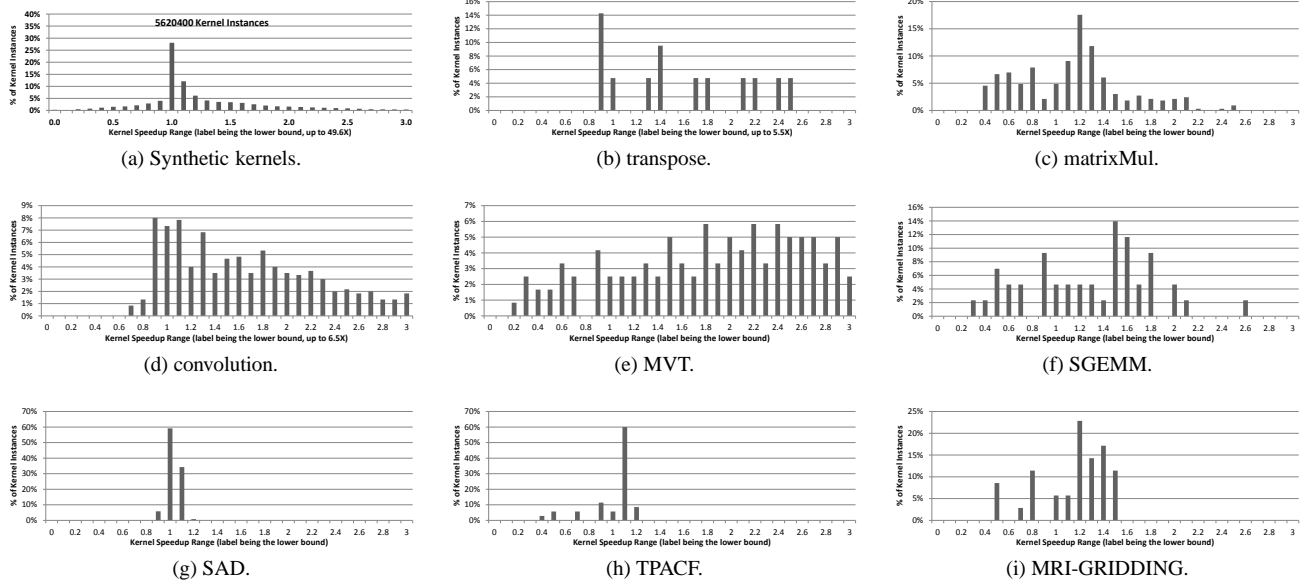


Figure 1. Histograms of kernel speedup brought by the local memory optimization.

penalty. These non-coalesced accesses can be eliminated by first copying a batch of the columns to the local memory in a coalesced manner. Workitems access data from the local memory, still one column at a time, but with no performance penalty. This is because, unlike global memory, local memory does not suffer from non-coalescing.

The copying of an array region from global memory to local memory is performed *cooperatively* by all workitems in a workgroup. The region is divided into a sequence of row segments, each having a width of a single DRAM transaction and is aligned to the transaction boundary [11]. These segments are cyclically distributed among *warps* in the workgroup. Elements in each segment are accessed by workitems in the designated warp in a fully coalesced manner, resulting in a single DRAM transaction. Overall, all global memory accesses made during the copying process are fully coalesced.

3. Impact of Using Local Memory

While the use of local memory reduces the number of DRAM transactions, it may not always improve the performance of the kernel as a whole. First, the optimization introduces the overhead of copying array regions from the global memory to the local memory. Second, it may reduce the level of parallelism, i.e., the number of threads that can concurrently execute on a GPU multi-processor, due to additional resource usage. A reduction in parallelism can hurt kernel performance in a holistic manner, potentially exposing latencies of *all* memory accesses. The extent to which this happens is a function of how many memory accesses exist in the kernel as well as the amount of computation in the kernel that may help *hide* this performance penalty.

Therefore, the performance impact of using local memory depends on various kernel characteristics:

- Amount of data reuse of the array region copied to the local memory
- Degree of memory non-coalescing of array accesses
- Usage of registers and local memory by the optimization, which can reduce parallelism
- Memory accesses and computation in the unoptimized kernel, which influences the performance impact of any parallelism drop

In order to assess the performance impact of these factors, we synthetically generate a large number of kernels with varying values of the characteristics listed above, and evaluated eight real-world benchmarks with varying launch configurations and other kernel parameters (Section 5). For each kernel, we empirically determine the kernel speedup of the local memory optimization, as the ratio of the execution time of the original kernel over that of the optimized kernel. Figure 1 shows the histograms of the resulting speedup values for both synthetic and real-world kernels. It confirms that the use of local memory is not always beneficial and its performance impact is non-trivial to determine.

4. ML-based Auto-tuning Framework

Given a kernel (with its launch configuration) containing accesses to an array that may exhibit data reuse or memory non-coalescing, our framework decides if the local memory optimization would improve kernel performance, by caching the smallest array region that covers these accesses in the local memory.




Figure 2. Overview of the machine-learning-based auto-tuning framework for the use of local memory.

The framework consists of two phases, as shown in Figure 2. In the first phase (the left part of the figure), a model is built using a machine learning algorithm, based on performance data of a set of *training OpenCL kernels* with and without the local memory optimization. The model correlates 1) characteristics (or *features*) extracted from a kernel, e.g., the degree of data reuse and the presence of memory non-coalescing and 2) optimization configuration, i.e., the shape of the array region to be cache, with 3) the benefit of the optimization, e.g., kernel speedup.

In the second phase (the right part of Figure 2), given a new (non-training) kernel along with the candidate array accesses, we extract kernel features as required by the model, determine the shape of the array region to be cached in local memory, and apply the model to predict if the optimization is beneficial.

Since a large number of training kernels is required for machine learning to work well, we opt to synthetically generate them instead of using real-world kernels. In the rest of the section, we discuss the design of the synthetic kernels and the machine learning model.

4.1 Synthetic Benchmarks

We design the synthetic kernels in the form of a single kernel template with a number of compile-time and run-time parameters. These parameters are “knobs” to alter kernel characteristics that may influence the benefit of using local memory (Section 3).

The kernel template is shown in Figure 3. It processes data from a 2D input array *in* and writes the result to a 2D output array *out*. We call the amount of work that produces an output array element a *work unit*, shown between lines 14 and 33. It contains two nested loops followed by a code segment, which we call *epilogue*. The loop nest encloses one or more accesses to array *in*, interleaved with computation (fused-multiply-add operations) and accesses to an auxiliary input array *in2*, a third kernel parameter. The epilogue also contains computation and accesses to *in2*, and ends with a write to array *out*. Accesses to array *in* are those to be

possibly cached in local memory, so we also refer to array *in* as the *target array*.

```

1  __kernel void kmain(
2      __global float *in ,
3      __global float *out ,
4      __global float *in2 ,
5      __local float *lmem)
6  {
7      int wg_x = get_group_id(0);
8      int wg_y = get_group_id(1);
9      int wi_x = get_local_id(0);
10     int wi_y = get_local_id(1);
11     ...
12     for iter_x = 0..(NUM_WUS_X-1) {
13         for iter_y = 0..(NUM_WUS_Y-1) {
14             int wu_x, wu_y;
15             (wu_x, wu_y) = func(wg_x, wg_y,
16                               wi_x, wi_y,
17                               iter_x, iter_y);
18             // This is where to cooperatively load
19             // a region of <in> to the local memory.
20             // barrier(...);
21             for i = 0..(N-1) {
22                 for j = 0..(M-1) {
23                     int idx_o = fo(wu_x, wu_y, i, j);
24                     int idx_i = fi(wu_x, wu_y, i, j);
25                     ... = in[idx_o + CO_1][idx_i + CI_1];
26                     ... // context (inner loop body)
27                     ... = in[idx_o + CO_k][idx_i + CI_k];
28                     ... // context (inner loop body)
29                 }
30             }
31             // barrier(...);
32             ... // context (epilogue)
33             out[...] = ...;
34         }
35     }
36 }

```

Figure 3. Synthetic kernel template.

Accesses to array *in* in the inner loop body are centered around a *home coordinate* (*idx_o*, *idx_i*), with different constant offsets in each dimension: *CO_1*, ..., *CO_k*, *CI_1*, ..., *CI_k* (lines 25 and 27). The home coordinate is a linear function of the current work unit coordinate (*wu_x*, *wu_y*) and the loop iterators *i* and *j*, as specified by *fo* and *fi* at lines 23 and 24.

The 2D grid of work units is distributed across a 2D space of workgroups in a blocked manner, and is further distributed across a 2D space of workitems in a cyclic manner. The geometry of array *out* and the launch configuration collectively determine the number of work units each workitem processes (*NUM_WUS_X* and *NUM_WUS_Y* at lines 12 and 13). The current work unit coordinate (*wu_x*, *wu_y*) is computed based on work group ID (*wg_x*, *wg_y*), work item ID (*wi_x*, *wi_y*) and work unit ID (*iter_x*, *iter_y*), at line 15.

The kernel template provides the flexibility to vary those characteristics that may affect the benefit of caching *in* data in the local memory. First, the data access pattern to the




Figure 4. Home access pattern visualization.

target array `in` is configurable. It is collectively defined by 1) the function tuple (`fo`, `fi`) that determines the *home access pattern*, and 2) the offsets `CO`'s and `CI`'s that determine the *stencil pattern* (within each workitem). We design 7 function tuples, shown in Figure 4, that correspond to regular access patterns with potentially different degrees of data reuse and memory non-coalescing. In each diagram, an arrow indicates the order of home coordinates (of array `in`) a workitem accesses as it goes through the iterations of loop `i` and `j` (line 21 and 22 of Figure 3). Each arrow is associated with a label that indicates what workitems of a workgroup make such accesses. For example, `WI(1,*)` refers to all workitems with `wi_x = 1`. Combined with the stencil pattern, this label reflects the amount of data reuse. The entire grey region shown in each diagram, extended with apron regions that cover neighbouring accesses, corresponds to the region of `in` to be cached in the local memory, reflecting the amount of resources consumed. We use three common stencil patterns: rectangular, diamond and star, shown in Figure 5. The target array `in` is padded to ensure no out-of-bound accesses.




Figure 5. Stencil patterns of target array accesses. The element of home coordinate is labeled with “H”.

In addition to the data access pattern, the amount of computation and memory accesses in the inner loop body and the epilogue are also configurable. They collectively model *contextual* kernel computation and memory accesses that may affect the optimization’s benefit with a drop in parallelism. Further, they enable variation of kernel register usage.

Table 1 summarizes the list of 13 parameters the kernel template provides. Run-time parameters are shown in *italics*; the rest are compile-time parameters.

4.2 Model Design

We build a machine-learning-based model that predicts kernel speedup brought by the local memory optimization. It

takes 18 inputs (or features) and outputs a single real number that reflects the kernel speedup. These inputs are:

1. Degree of data reuse exhibited by the home access, i.e., the average number of workitems in a workgroup where this access refers to the same array element.
2. Amount of local memory used by each workgroup for the optimization.
3. Degree of non-coalescing exhibited by the home access (in the unoptimized kernel), i.e., the average number of memory transactions induced by a warp.
4. Number of accesses to the target array.
5. Minimum and maximum offsets to the home coordinate of target array accesses (2 parameters in each dimension).
6. Number of computation operations in the inner loop body and the epilogue (2 parameters).
7. Number of contextual memory accesses (i.e., not made to the target array) in the inner loop body and the epilogue, and whether each is coalesced or not (4 parameters).
8. Number of registers used per thread (in the unoptimized kernel).
9. Grid size and workgroup size (2 parameters).
10. Number of work units each workitem processes, equivalent to `NUM_WUS_X * NUM_WUS_Y` in the synthetic kernel template.

It is easy to extract the above features from a synthetic kernel, because they can be directly mapped to the parameters of the kernel template. For example, features #1-#5 are computed from the template parameters that define the data access pattern. Features #6 and #7 are one-to-one correspondent to the template parameters that define the kernel context. Currently features are extracted automatically from the synthetic kernels *when they are generated*.

To extract features from a real-world application, we must first map the kernel structure to that of the synthetic kernel template, by identifying the boundary of a work unit (Section 4.1). We currently extract features from real-world applications *manually*. However, we believe that a compiler would be able to extract them automatically once the work unit boundary is identified.

Table 1. Parameters of the synthetic kernel template.

Category	Parameter	Description
Global	IN_H, IN_W	Height and width of the target array <code>in</code>
Data Access Pattern	HOME_ACCESS_PATTERN	One of the seven shown in Figure 4
	N, M	Trip-counts of loop <code>i</code> and <code>j</code>
	STENCIL_PATTERN	One of rectangular, diamond and star
	STENCIL_RADIUS	Radius of the selected stencil pattern
Kernel Context	NUM_COMP_ILB/EP	# of computation in the inner loop body and epilogue
	NUM_COAL_ACCESSES_ILB/EP	# of coalesced accesses to array <code>in2</code> in the inner loop body and epilogue
	NUM_UNCOAL_ACCESSES_ILB/EP	# of non-coalesced accesses to array <code>in2</code> in the inner loop body and epilogue

5. Evaluation

We generate a total of 9600 synthetic kernels from the template, by varying the parameters in Table 1. We evaluate each synthetic kernel with a number of launch configurations, resulting in a total of 5.6 million *kernel instances*. We run each kernel instance with and without the local memory optimization.

We select the values of the kernel template parameters in two steps. First, we randomly sample 100 tuples from all compile-time parameters except HOME_ACCESS_PATTERN, with the resulting value distributions listed in Table 2. Second, for each tuple, we enumerate all 7 home access patterns. For each pattern, we enumerate a set of 4 values for N and 4 values for M , that we perceive as common. The value set for N is 8, 16, 32, 64 for home access patterns `xy-reuse` and `x/y-reuse-row`, and 1, 2, 4, 8 for others. The value set for M is 8, 16, 32, 64 for home access patterns `xy-reuse` and `x/y-reuse-col`, and 1, 2, 4, 8 for others. The target array shape ($IN_H \times IN_W$) is fixed at 2048×2048 .

Table 2. Compile-time parameter value distribution for synthetic kernels.

Parameter	Value Range (Average)
STENCIL_PATTERN	All three
STENCIL_RADIUS	0 – 2
NUM_COMP_ILB	5 – 44 (19)
NUM_COMP_EP	1 – 48 (23)
NUM_COAL_ACCESSES_ILB	0 – 13 (3)
NUM_COAL_ACCESSES_EP	0 – 13 (5)
NUM_UNCOAL_ACCESSES_ILB	0 – 4 (0.8)
NUM_UNCOAL_ACCESSES_EP	0 – 4 (0.8)

For each synthetic kernel, we sweep through: 1) all valid 2D grid geometries with individual dimensions restricted to powers of 2 and the total size no less than 512, and 2) all valid 2D workgroup geometries with individual dimensions restricted to powers of 2 and the total size no more than 1024.

In addition to synthetic kernels, we also look at eight real-world benchmarks, summarized in Table 3. For each one, we vary kernel parameters such as launch configurations and tiling factors, resulting in multiple kernel instances.

We collect the execution time of all kernel instances (both synthetic and real) on NVIDIA Tesla M2090 with 6GB of memory, housed in a system with an Intel Xeon E5-2620 CPU and 64GB of memory, running CUDA 5.0 on CentOS 6.4. We measure the execution time of the kernel only.

Figure 1a showed the distribution of synthetic kernel speedup values brought by the local memory optimization. Figure 1b- 1i showed the distribution of real-world kernel speedup values. We make two observations from the figures. First, the use of local memory is not always beneficial for both synthetic and real-world kernels. Second, the speedup distributions have different shapes across the synthetic and real-world kernels. This demonstrates the need for auto-tuning the use of local memory.

5.1 Model Training and Evaluation

We train a model using the performance data of randomly selected 560K synthetic kernel instances (10% of the total). We build the model using Random Forest (RF) [2], from Weka 3.7.10 [7], configured with 20 trees (of unlimited depth) and 4 attributes per tree node.


We evaluate the model’s accuracy by applying it to the remaining synthetic kernel instances and all real-world kernel instances. We use two accuracy metrics. The first is *count-based accuracy*, defined as the percentage of kernel instances where the decision of whether or not to use local memory, predicted by the model, matches the *oracle* decision based on actual kernel performance data. Effectively this metric assigns to each kernel instance a score of 1 when the model predicts correctly and 0 otherwise, and computes the average across all kernel instances. Note that, when the model mis-predicts, the metric does not take into account the performance loss it incurs. Hence we introduce the second metric, *penalty-weighted accuracy*, which extends the first metric by assigning a score equal to the performance ratio (a value between 0 and 1), instead of 0, when the model mis-predicts. Effectively this metric measures the percentage of kernel performance achieved using the model-predicted decision, over that achieved by the oracle decision, averaged across all kernel instances.

Figure 5.1 shows the model accuracy, with both metrics, for synthetic and real-world kernels. For penalty-

Table 3. Real-world benchmarks.

Benchmark	Suite	Description	LOC	# of Kernel Instances
transpose	NVIDIA SDK	Matrix transpose	6	21
matrixMul		Matrix multiply ($C = A \times B$)	9	330
convolution		2D separable convolution	10	600
MVT	Polybench	Matrix vector multiply	9	120
SGEMM	Parboil	$C = \alpha \times A \times B + \beta \times C$	10	48
SAD		Computes Sum-of-Absolute-Differences between pairs of image blocks; used in motion estimation algorithm in H.264	94	517
TPACF		Computes the angular correlation function for a data set of astronomical bodies	129	35
MRI-GRIDDING		Computes a regular grid of data representing an MR scan by weighted interpolation of actual acquired data points	126	35

weighted accuracy, we also show the range (min-max) of per-kernel-instance scores using error bars. Overall, the trained model achieves 86% count-based accuracy and nearly 95% penalty-weighted accuracy on the remaining synthetic kernel instances. However, a 30% minimum score indicates that the model does mis-predict on a small percentage of kernel instances with high performance penalty. For real-world benchmarks, the model is able to achieve nearly 95% penalty-weighted accuracy, although the count-based accuracy drops noticeably for SAD, TPACF and MRI-GRIDDING. This shows that the model trained with a large number of synthetic kernels can achieve high accuracy consistently across a variety of kernels.

**Figure 6.** Accuracy of machine-learning-based models.

6. Related Work

There has been growing interest in the use of machine learning to auto-tune the performance of GPU applications [1, 6, 8–10]. For example, Magni et al. [10] explore the use of neural network for auto-tuning thread coarsening, and Grewe et al. [6] use a decision tree to decide if an OpenCL kernel should be executed on the CPU or the GPU. In contrast, we focus on auto-tuning the use of local memory.

There is work that explored auto-tuning of the use of local memory, but focused on the use of analytical modeling and empirical search [12]. In contrast, we build machine learning

models, which have the potential to be more accurate than analytical approaches.

Finally, there is a large body of work that treats auto-tuning for platforms other than GPUs, including multi-cores [5, 16] and single-core processors [3, 4, 13]. In contrast, we focus on GPUs.

7. Conclusion and Future Work

In this paper we described a machine learning model for use in automatic performance tuning of local memory usage on GPUs. We have shown that the optimization is not always beneficial, depending on application characteristics. We train a Random Forest model with a large number of synthetic benchmarks and predict whether local memory should be used or not for a single array access in both synthetic and real benchmarks. We have shown that the penalty-weighted prediction accuracy is nearly 95%.

This work can be extended in various directions. First, the use of this model in practice demands a compiler framework that automatically applies the local memory optimization and extracts kernel features. Second, the model can be extended to predict the *usage* of local memory when multiple arrays compete for local memory resources. Third, the prediction accuracy can be evaluated for a larger set of real-world benchmarks. Fourth, the quality of the machine learning model using synthetic benchmarks, as opposed to real benchmarks, can be evaluated. Finally, other machine learning models (e.g., Support Vector Machines) can be evaluated.

Acknowledgments

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC) and Qualcomm, Inc., and by an NVIDIA Graduate Fellowship.

References

- [1] J. Bergstra, N. Pinto, and D. Cox. Machine learning for predictive auto-tuning with boosted regression trees. In *Proc.*

- of *Innovative Parallel Computing (InPar)*, pages 1–9, May 2012.
- [2] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, Oct. 2001.
 - [3] J. Cavazos and J. E. B. Moss. Inducing heuristics to decide whether to schedule. In *Proc. of PLDI*, pages 183–194, 2004.
 - [4] G. Fursin et al. Milepost GCC: Machine Learning Enabled Self-tuning Compiler. *International Journal of Parallel Programming*, 39(3), 2011.
 - [5] A. Ganapathi, K. Datta, A. Fox, and D. Patterson. A case for machine learning to optimize multicore performance. In *Proc. of HotPar*, pages 1–1, 2009.
 - [6] D. Grewe, Z. Wang, and M. O’Boyle. Portable mapping of data parallel programs to opencl for heterogeneous systems. In *Proc. of CGO*, pages 1–10, Feb 2013.
 - [7] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, Nov. 2009. ISSN 1931-0145.
 - [8] Y. Liu, E. Zhang, and X. Shen. A cross-input adaptive framework for GPU program optimizations. In *Proc. of IPDPS*, pages 1–10, 2009.
 - [9] A. Magni, D. Grewe, and N. Johnson. Input-aware auto-tuning for directive-based gpu programming. In *Proc. of GPGPU*, pages 66–75, 2013.
 - [10] A. Magni, C. Dubach, and M. O’Boyle. Automatic optimization of thread-coarsening for graphics processors. In *Proc. of PACT*, pages 455–466, 2014.
 - [11] NVIDIA. NVIDIA CUDA C Programming Guide v5.0, 2013.
 - [12] S. Ryoo et al. Program optimization carving for GPU computing. *J.Parallel Distrib.Comput.*, 68(10):1389–1401, 2008.
 - [13] M. Stephenson, S. Amarasinghe, M. C. Martin, and U.-M. O’Reilly. Meta optimization: Improving compiler heuristics with machine learning. In *Proc. of PLDI*, Jun 2003.
 - [14] The Khronos Group. Open Computing Language (OpenCL). <http://www.khronos.org/opencl/>.
 - [15] S.-Z. Ueng et al. CUDA-Lite: Reducing GPU Programming Complexity. In *Proc. of LCPC*, pages 1–15, 2008.
 - [16] Z. Wang and M. F. O’Boyle. Partitioning streaming parallelism for multi-cores: a machine learning based approach. In *Proc. of PACT*, pages 307–318, 2010.